ABSTRACT OF THE DISCLOSURE

[0063] A method for assessing how long continuously operating software systems can be expected to remain executing in a safe and/or reliable manner before anomalous conditions will ultimately lead to failure. For safety-critical applications the method can provide a safe upper bound on the time between rebooting. Also disclosed is an empirical technique for determining which portions of the state, if corrupted, create the greatest risks to safe and/or reliable, continual execution of the software. Armed with this information, developers, testers, and certifiers can create justifiable plans for the frequency with which the software should be rebooted. Further, they can customize and embed internal self-tests into those portions of the state found to have the greatest risks to safe and/or reliable, continual execution of the software. These self-tests can also warn when failures are likely to occur well in advance of the failures, so that the software may be safely rejuvenated to avert undesired or catastrophic conclusions